

(19) World Intellectual Property Organization
International Bureau(43) International Publication Date
3 April 2003 (03.04.2003)

PCT

(10) International Publication Number
WO 03/027908 A2

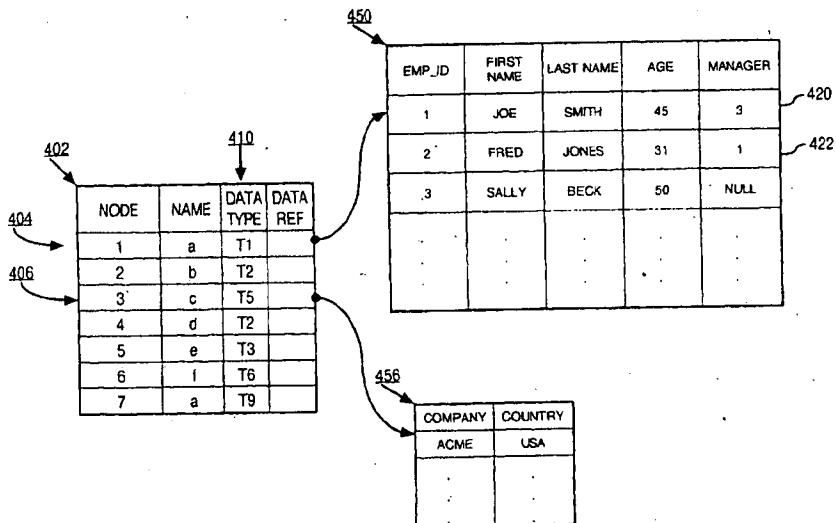
- (51) International Patent Classification⁷: G06F 17/30
- (21) International Application Number: PCT/US02/31168
- (22) International Filing Date: 27 September 2002 (27.09.2002)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
60/326,052 28 September 2001 (28.09.2001) US
60/378,800 7 May 2002 (07.05.2002) US
- (71) Applicant: ORACLE INTERNATIONAL CORPORATION [US/US]; 500 Oracle Parkway, M/S 5op7, Redwood Shores, CA 94065 (US).
- (72) Inventors: AGARWAL, Nipun; 4768 Cheeney Street, Santa Clara, CA 95054 (US). SEDLAR, Eric; 4270 Cesar Chavez Street, San Francisco, CA 94131 (US). MURTHY, Ravi; 33227 Jamie Circle, Fremont, CA 94555 (US). JAIN, Namit; 1103 Foster City Blvd, Foster City, CA 94404 (US).
- (74) Agents: HICKMAN, Brian; 1600 Willow Street, San Jose, CA 95125 et al. (US).
- (81) Designated States (national): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VC, VN, YU, ZA, ZM, ZW.
- (84) Designated States (regional): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— without international search report and to be republished upon receipt of that report

[Continued on next page]

(54) Title: PROVIDING A CONSISTENT HIERARCHICAL ABSTRACTION OF RELATIONAL DATA



(57) Abstract: A hierarchy may be explicitly or implicitly reflected in existing data maintained external to hierarchy structures. Such hierarchies are referred to herein as a "pre-existing hierarchies". Techniques are described herein for capturing a pre-existing hierarchy in hierarchy structures. After a pre-existing hierarchy has been captured, there exist two independent sets of data that reflect the hierarchy: the external hierarchy definition and the internal hierarchy definition. Changing either hierarchy definition changes the hierarchy. Therefore, to maintain a consistent reflection of the hierarchy, the internal hierarchy definition must be modified in response to changes made to the external hierarchy definition, and the external hierarchy definition must be modified in response to changes made to the internal hierarchy definition. Various techniques are described for maintaining consistency between the two hierarchy definitions.

WO 03/027908 A2



For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

**PROVIDING A CONSISTENT HIERARCHICAL ABSTRACTION
OF RELATIONAL DATA**

RELATED APPLICATIONS

The present application claims priority from the following U.S. Provisional Patent Applications, the entire contents of which are incorporated herein by reference for all purposes:

U.S. Provisional Patent Application No. 60/326,052, filed on September 28, 2001, entitled "File Based Access Provided With a Database System," by Eric Sedlar and Viswanathan Krishnamurthy;

U.S. Provisional Patent Application No. 60/378,800, filed on May 7, 2002, entitled "SQL Access to Data that Provides a File System Abstraction," by Nipun Agarwal, Ravi Murthy, Eric Sedlar, Sivasankaran Chandrasekar, Fei Ge, Syam Pannala, Neema Jalali and Muralidhar Krishnaprasad.

The present application is also related to the following U.S. Patent Applications, the entire contents of which are incorporated herein by reference for all purposes:

U.S. Patent Application Serial No. _____, filed on the equal day herewith, entitled OPERATORS FOR ACCESSING HIERARCHICAL DATA IN A RELATIONAL SYSTEM, by Nipun Agarwal, Ravi Murthy, Eric Sedlar, Sivasankaran Chandrasekar and Fei Ge (Attorney Docket No. 50277-1975);

U.S. Patent Application Serial No. _____, filed on the equal day herewith, entitled MECHANISM FOR MAPPING XML SCHEMAS TO OBJECT-RELATIONAL DATABASE SYSTEMS, by Ravi Murthy, Muralidhar Krishnaprasad, Sivasankaran Chandrasekar, Eric Sedlar, Vishu Krishnamurthy and Nipun Agarwal (Attorney Docket No. 50277-1977);

U.S. Patent Application Serial No. _____, filed on the equal day herewith, entitled INDEXING TO EFFICIENTLY MANAGE VERSIONED DATA IN A DATABASE SYSTEM, by Nipun Agarwal, Eric Sedlar and Ravi Murthy (Attorney Docket No. 50277-1978);

U.S. Patent Application Serial No. _____, filed on the equal day herewith, entitled MECHANISMS FOR STORING CONTENT AND PROPERTIES OF HIERARCHICALLY ORGANIZED RESOURCES, by Ravi Murthy, Eric Sedlar, Nipun Agarwal, and Neema Jalali (Attorney Docket No. 50277-1979);

U.S. Patent Application Serial No. _____, filed on the equal day herewith, entitled MECHANISM FOR UNIFORM ACCESS CONTROL IN A DATABASE SYSTEM, by Ravi Murthy, Eric Sedlar, Nipun Agarwal, Sam Idicula, and Nicolas Montoya (Attorney Docket No. 50277-1980);

U.S. Patent Application Serial No. _____, filed on the equal day herewith, entitled LOADABLE UNITS FOR LAZY MANIFESTATION OF XML DOCUMENTS by Syam Pannala, Eric Sedlar, Bhushan Khaladkar, Ravi Murthy, Sivasankaran Chandrasekar, and Nipun Agarwal (Attorney Docket No. 50277-1981);

U.S. Patent Application Serial No. _____, filed on the equal day herewith, entitled MECHANISM TO EFFICIENTLY INDEX STRUCTURED DATA THAT PROVIDES HIERARCHICAL ACCESS IN A RELATIONAL DATABASE SYSTEM, by Neema Jalali, Eric Sedlar, Nipun Agarwal, and Ravi Murthy (Attorney Docket No. 50277-1982).

FIELD OF THE INVENTION

The present invention relates to storing, in a relational database, information that is organized according to a hierarchy and, more specifically, to techniques for managing the metadata that captures the hierarchy.

BACKGROUND OF THE INVENTION

Humans tend to organize information in categories. The categories in which information is organized are themselves typically organized relative to each other in some form of hierarchy. For example, an individual animal belongs to a species, the species belongs to a genus, the genus belongs to a family, the family belongs to an order, and the order belongs to a class.

With the advent of computer systems, techniques for storing electronic information have been developed that largely reflected this human desire for hierarchical organization. Conventional computer file systems, for example, are typically implemented using hierarchy-based organization principles. Specifically, a typical file system has directories arranged in a hierarchy, and documents stored in the directories. Ideally, the hierarchical relationships between the directories reflect some intuitive relationship between the meanings that have been assigned to the directories. Similarly, it is ideal for each document to be stored in a directory based on some intuitive relationship

between the contents of the document and the meaning assigned to the directory in which the document is stored.

Fig. 1 shows an example of a typical file system. The illustrated file system includes numerous directories arranged in a hierarchy. Two documents 118 and 122 are stored in the directories. Specifically, documents 118 and 122, both of which are entitled "Example.doc", are respectively stored in directories 116 and 124, which are respectively entitled "Word" and "App4".

In the directory hierarchy, directory 116 is a child of directory 114 entitled "Windows", and directory 114 is a child of directory 110. Similarly, directory 124 is a child of directory 126 entitled "VMS", and directory 126 is a child of directory 110. Directory 110 is referred to as the "root" directory because it is the directory from which all other directories descend. In many systems, the symbol "/" is used to refer to the root directory.

When electronic information is organized in a hierarchy, each item of information may be located by following a "path" through the hierarchy to the entity that contains the item. Within a hierarchical file system, the path to an item begins at the root directory and proceeds down the hierarchy of directories to eventually arrive at the directory that contains the item of interest. For example, the path to file 118 consists of directories 110, 114 and 116, in that order.

Hierarchical storage systems often allow different items to have the same name. For example, in the file system shown in Figure 1, both of the documents 118 and 122 are entitled "Example.doc". Consequently, to unambiguously identify a given document, more than just the name of the document is required.

A convenient way to identify and locate a specific item of information stored in a hierarchical storage system is through the use of a "pathname". A pathname is a concise way of uniquely identifying an item based on the path through the hierarchy to the item. A pathname is composed of a sequence of names. In the context of a file system, each name in the sequence of names is a "filename". The term "filename" refers to both the names of directories and the names of documents, since both directories and documents are considered to be "files".

Within a file system, the sequence of filenames in a given pathname begins with the name of the root directory, includes the names of all directories along the path from the root directory to the item of interest, and terminates in the name of the item of interest. Typically, the list of directories to traverse is concatenated together, with some

kind of separator punctuation (e.g., '/', '\', or ';') to make a pathname. Thus, the pathname for document 118 is /Windows/Word/Example.doc, while the pathname for document 122 is /VMS/App4/Example.doc.

The relationship between directories (files) and their contained content varies significantly between different types of hierarchically organized systems. One model, employed by various implementations, such as Windows and DOS file systems, requires each file to have exactly one parent, forming a tree. In a more complicated model, the hierarchy takes the form of a directed graph, where files can have multiple parents, as in the UNIX file system in which hard links are used.

In contrast to hierarchical approaches to organizing electronic information, a relational database stores information in tables comprised of rows and columns. Each row is identified by a unique RowID. Each column represents an attribute of a record, and each row represents a particular record. Data is retrieved from the database by submitting queries to a database management system (DBMS) that manages the database. The queries must conform to the database language supported by the database management system. Structured Query Language (SQL) is an example of a database language supported by many existing database management systems.

Each type of storage system has advantages and limitations. A hierarchically organized storage system is simple, intuitive, and easy to implement, and is a standard model used by most application programs. Unfortunately, the simplicity of the hierarchical organization does not provide the support required for complex data retrieval operations. For example, the contents of every directory may have to be inspected to retrieve all documents created on a particular day that have a particular filename. Since all directories must be searched, the hierarchical organization does nothing to facilitate the retrieval process.

A relational database system is well suited for storing large amounts of information and for accessing data in a very flexible manner. Relative to hierarchically organized systems, data that matches even complex search criteria may be easily and efficiently retrieved from a relational database system. However, the process of formulating and submitting queries to a database server is less intuitive than merely traversing a hierarchy of directories, and is beyond the technical comfort level of many computer users.

In the past, hierarchically organized systems and relationally organized systems have been implemented in different ways that were not compatible. With some additional

processing, however, a relationally organized system can emulate a hierarchically organized system. This type of emulation is especially desirable when the storage capability and flexibility of a relational system is needed, but the intuitiveness and ubiquity of the hierarchical system is desired.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention is illustrated by way of example, and not by way of limitation, in the figures of the accompanying drawings and in which like reference numerals refer to similar elements and in which:

FIG. 1 is a block diagram showing a hierarchical file system;

FIG. 2 is a block diagram showing a hierarchy of nodes that are associated with resources;

FIG. 3 is a block diagram of relational tables that can be used to capture the hierarchy illustrated in FIG. 2, according to an embodiment of the invention;

FIG. 4 is a block diagram illustrating the separation of hierarchy structures and content structures according to an embodiment of the invention;

FIG. 5 is a block diagram of a relational table in which non-leaf nodes of an information hierarchy are indicated by values stored in rows that correspond to the leaf nodes;

FIG. 6 is a block diagram of the hierarchy implicit in the data stored in the table of FIG. 5; and

FIG. 7 is a block diagram of a system on which embodiments of the invention may be implemented.

DETAILED DESCRIPTION OF EMBODIMENTS OF THE INVENTION

A method and system are described for storing resources that belong to an information hierarchy into structures within a relational database system. In the following description, for the purposes of explanation, numerous specific details are set forth in order to provide a thorough understanding of the present invention. It will be apparent, however, that the present invention may be practiced without these specific details. In other instances, well-known structures and devices are shown in block diagram form in order to avoid unnecessarily obscuring the present invention.

FUNCTIONAL OVERVIEW

Various techniques are provided for facilitating the management of hierarchical data within a relational database system. According to one embodiment, the storage structures used to store the content of the resources that belong to a hierarchy (the "content structures") are separate from the storage structures used to store data that captures the information about the hierarchy (the "hierarchy structures").

In some situations, a hierarchy may already be explicitly or implicitly reflected in existing data maintained external to the hierarchy structures. Such hierarchies are referred to herein as a "pre-existing hierarchies". In these situations, it is desirable to capture the pre-existing hierarchy in the hierarchy structures so that the hierarchy structures may be used as an alternative access path to the resources. For example, the hierarchy structures may be used to access the resources based on path information. Further, even when the resources are accessed by issuing queries directly against the content structures, the access operations gain the benefits of any features and access structures associated with the hierarchy structures, such as any security mechanism built into the hierarchy structures and any indexes built to increase the performance of operations that access resources based on their position within the hierarchy.

Each pre-existing hierarchy is associated with data that (1) resides external to the hierarchy structures, and (2) reflects the pre-existing hierarchy. Such data is referred to herein as an "external hierarchy definition". To capture a pre-existing hierarchy in the hierarchy structures, data (referred to herein as the "internal hierarchy definition") must be added to the hierarchy structure based on the external hierarchy definition.

After a pre-existing hierarchy has been captured, there exist two independent sets of data that reflect the hierarchy: the external hierarchy definition and the internal hierarchy definition. Changing either hierarchy definition changes the hierarchy. Therefore, to maintain a consistent reflection of the hierarchy, the internal hierarchy definition must be modified in response to changes made to the external hierarchy definition, and the external hierarchy definition must be modified in response to changes made to the internal hierarchy definition. Various techniques are described herein for maintaining consistency between the two hierarchy definitions.

STORING HIERARCHICAL DATA IN A DATABASE SYSTEM

FIG. 2 is a block diagram that illustrates a hierarchy 200 that is used in examples that are provided herein to facilitate an understanding of embodiments of the invention.

Hierarchy 200 includes eight nodes. The highest node in the hierarchy is referred to as the "root" node. The nodes at the end of each branch in the hierarchy are "leaf" nodes. The nodes between the root node and the leaf nodes are "intermediate" nodes. In the illustrated hierarchy, nodes 1, 2, and 3 are intermediate nodes, and nodes 4, 5, 6, and 7 are leaf nodes.

In an information hierarchy, the nodes correspond to information. Typically, the piece of information associated with each node will have some form of name, and some type of content. For example, in a hierarchy that corresponds to a hierarchical file system, the nodes will typically correspond to files (where a "folder" or "directory" is one type of file). Each such file will have a name, and some form of contents.

In many situations, the names that are associated with the nodes in a hierarchy need not be unique. In the example shown in FIG. 2, both node 1 and node 7 have the name "a". Though they have the same name, they are separate and distinct nodes, which occupy different positions in the hierarchy and may be associated with completely different content.

ONE-SIZE-FITS-ALL

FIG. 3 is a block diagram of two tables (resource table 302 and link table 350) that may be used to represent hierarchy 200 in a relational database system. Resource table 302 includes one row for each node in the hierarchy. The row for a given node includes the name associated with the given node and the data associated with the given node. For example, row 304 corresponds to node 1, and contains the name "a" and data 306 associated with node 1. Resource table 302 is an example of a one-size-fits-all approach in that the data for all resources in the hierarchy are stored in the same structure, regardless of the data type of those resources. In the example of resource table 302, the data structure used to store the content of the resources is the single column 308. The data type for such a column could be, for example, a LOB (Large Object) type, such as a Binary LOB (BLOB) or a Character LOB (CLOB).

Link table 350 includes one row for each of the parent-child relationships in hierarchy 200. A parent column 378 holds the value that specifies the parent in a parent-child relationship, and a child column 380 holds the value that specifies the child of the parent-child relationship. For example, row 352 indicates that the root node is the parent of node 1. Similarly, rows 354 and 356 respectively indicate that node 1 is the parent of nodes 2 and 3.

SEPARATING HIERARCHICAL STRUCTURES FROM CONTENT STRUCTURES

Tables 302 and 350 capture all the information of the hierarchy 200 in relational format. However, it may not be possible to fully utilize the power of the relational database system if the content of the resources associated with hierarchy 200 are stored using the one-size-fits-all approach represented by resource table 302. Specifically, resource table 302 includes a single column for storing the content of the resources associated with the nodes in the hierarchy 200. This storage format may be sufficient when the resource associated with a node is a simple data type, such as a string or integer. However, in many situations the resources associated with the nodes in the hierarchy will be complex data types, such as complex records or XML documents. In such situations, the storage of the resource content in a single column inhibits the user's ability to run complex searches against the content. Even when the relational database system provides support for searching within a single column for values of subcomponents of complex resources, the performance of such searches is impaired.

According to one embodiment of the invention, the problems associated with storing and searching the content of complex resources that belong to an information hierarchy are overcome by storing the content of the resources associated with the nodes separately from the tables that capture the hierarchy information.

FIG. 4, for example, shows a resource table 402 similar to resource table 302. However, resource table 402 differs from resource table 302 in that resource table 402 no longer stores the content of each resource in a column of the row for that resource. Rather, the DATA column of resource table 302 is replaced in resource table 402 with two columns DATA TYPE and DATA REF.

For each row of resource table 402, the DATA TYPE column stores data that indicates the type of resource associated with the node identified in the row. For example, row 404 corresponds to node 1 of hierarchy 200. The DATA TYPE column of row 404 indicates that the resource associated with node 1 is of type T1. Similarly, row 406 corresponds to node 3 and the DATA TYPE column of row 406 indicates that the resource associated with node 3 is of type T5.

For each row of resource table 402, the DATA REF column stores a reference to the resource associated with the node identified in the row. As used herein, the term "reference" generally refers to any data used within a database system to locate other data. For example, the DATA REF column can store pointers, or more complex information

such as the "ref" data type supported by Oracle database systems. The present invention is not limited to any particular type of data for locating the resources associated with nodes in a hierarchy.

Various benefits are realized by storing the content of resources separately from the hierarchical structures that are used to capture and reflect the hierarchical relationships between the resources. For example, because the resources are stored separately from resource table 402, the content structures in which the resources are stored may be defined based on the nature of the resources. For example, the resource associated with node 1 is an employee record, and is stored in a table 450 that has columns for each of the fields of an employee record. On the other hand, the resource associated with node 3 is a company record, and is stored in a table 456 that has columns for each of the fields of a company record.

Also, because the hierarchy structures (e.g. tables 402 and 350) are separate from the content structures (e.g. tables 450 and 456), a hierarchy may be established for existing relational data with minimal effect on the storage of that data, and without requiring modification to the existing queries used to access that data. For example, the employee table 450 may have existed long before the decision to arrange information into the hierarchy 200. Also prior to that decision, many queries and views may have been defined for accessing table 450. Because tables 350 and 402 can be created without affecting the existing structure of table 450, the hierarchical structures for a hierarchy 200 that includes resources stored in table 450 may be established without the need to modify table 450 or the existing queries and views that access table 450.

According to one embodiment, resource table 402 may be implemented in a way that uses the one-size-fits-all approach for some resources, and contains references to the contents of other resources. For example, while the content of the resource associated with node 1 may be stored in table 450, the content of the resource associated with node 2 may continue to be stored in a single LOB column within table 402. The data type of a resource, as indicated in the DATA TYPE column, may be used by the database server as one factor in determining whether to store the resource content within a LOB column of the resource table 402 or in one or more separate tables.

PRE-EXISTING HIERARCHIES

Table 450 may initially exist completely independent of hierarchy structures, such as resource table 402 and link table 350. Thus, the hierarchy structures do not initially

capture any information about any hierarchy to which the data in table 450 may belong. However, information stored external to such hierarchy structures may explicitly or implicitly establish a hierarchy between the rows of table 450.

In the case of table 450, such a hierarchy is implicitly established by the information contained in the EMP_ID and MANAGER columns. Specifically, the data in the EMP_ID and MANAGER columns of table 450 establish an implicit "management hierarchy" in which: the employee associated with emp_id 1 is immediately below the employee associated with emp_id 3; the employee associated with emp_id 2 is immediately below the employee associated with emp_id 1; and the employee associated with emp_id 3 has no manager (and therefore is the root node of the management hierarchy).

In the present example, the management hierarchy is a pre-existing hierarchy and the data in stored in the EMP_ID and MANAGER columns constitute an external hierarchy definition. To use the hierarchy structures to access the data in table 450 based on the management hierarchy, information about the management hierarchy must be captured in the hierarchy tables. Various techniques for capturing information about a pre-existing hierarchy shall be described in greater detail hereafter.

CAPTURING A PRE-EXISTING HIERARCHY

According to one embodiment, capturing a pre-existing hierarchy involves: (1) storing in a resource table a row for each node in the pre-existing hierarchy, and (2) storing, in a hierarchy structure (such as a hierarchical index), information that captures the parent-child relationships between the nodes of the pre-existing hierarchy. Further, if the pre-existing hierarchy is being grafted into a larger hierarchy that is already reflected in the hierarchy structures, then metadata must be added to the hierarchy structures to establish the root node of the pre-existing hierarchy as a child of a node in the larger hierarchy.

Operations that perform the tasks of (1) storing in a resource table a row for each node in the pre-existing hierarchy, and (2) adding metadata that grafts the root node of the pre-existing hierarchy to a larger hierarchy are relatively straightforward. On the other hand, the technique used to capture the parent-child relationships between the nodes of the pre-existing hierarchy may vary based on the nature of the external hierarchy definition. For example, if an SQL relation defines the parent-child relationship, then the parent-child information may be captured using a SQL command that uses an appropriate

connect-by clause. Alternatively, if the parent-child relationship is implicit based on hierarchically related fields, then the parent-child information may be captured using an enumeration approach. Both the connect-by and enumeration approach will be described in greater detail hereafter.

AUTOMATING CAPTURE OF PRE-EXISTING HIERARCHIES

According to one embodiment, a database server is configured to perform pre-existing hierarchy capture operations automatically in response to receiving certain information about the pre-existing hierarchy. Specifically, according to one embodiment, a database server is configured to perform a capture operation automatically in response to receiving (1) information that identifies the content structure associated with the pre-existing hierarchy, and (2) connecting criteria that indicates criteria for establishing parent-child relationships between the resources contained in the specified content structure.

For the purpose of illustration, assume that it is desirable to capture the management hierarchy associated with table 450. Under such circumstances, the automated capture operation could be performed in response to receiving (1) data that identifies table 450, and (2) data that indicates that each given row of table 450 is a child of the row whose EMP_ID is equal to MANAGER specified in the given row. Techniques that may be used to automatically capture the pre-existing hierarchy based on this information shall be described hereafter.

CAPTURING PARENT-CHILD RELATIONSHIPS USING CONNECT BY

As mentioned above, the parent-child relationships of some pre-existing hierarchies may be indicated by an SQL relationship. For example, the management hierarchy of the records in table 450 is established by the data contained in the MANAGER column of table 450. In such cases, the parent-child relationships of the pre-existing hierarchy can be captured through the use of database commands that use an appropriate CONNECT BY clause.

For example, an automated capture operation of the management hierarchy is performed by (1) creating rows in the resource table 402 to represent nodes that correspond to the rows in table 450, and (2) generating metadata about the hierarchical relationships between those nodes by executing an SQL statement that includes a CONNECT BY clause. Each row in the resource table 402 that represents the node for a

particular row of table 450 has a reference to that particular row. For example, row 404 of resource table 402 represents the node for row 420 of table 450, and therefore includes a reference to row 420.

With respect to generating the path information associated with the pre-existing hierarchy, assume that table 450 is named 'EMP' and is defined to include the following columns: (emp_id integer, first_name varchar2(80), last_name varchar2(80), age integer, manager integer). Under these circumstances, a tree of entries (of variable depth, based on the data) can be generated using a CONNECT BY relationship such as 'START WITH manager IS NULL CONNECT BY PRIOR manager = emp_id'. In this situation, a user would typically also specify a column value to use as the name of the path element (e.g. 'first_name').

CAPTURING A PRE-EXISTING HIERARCHY USING ENUMERATION

In some cases, a pre-existing hierarchy may be implicitly defined by fields that have a hierarchical relationship with each other. For example, consider a relational table containing a list of products offered for sale. Such a table may, for example, be defined as follows:

```
CREATE TABLE PRODUCT_INFO
(
    manufacturer VARCHAR2(128),
    product_type VARCHAR2(48),
    product_name VARCHAR2(80),
    sku NUMBER
);
```

Such a table may be populated, for example, as illustrated by table 500 of FIG. 5. In this example, a hierarchy is implicit in the hierarchical relationship between the MANUFACTURER, PRODUCT_TYPE, PRODUCT_NAME and SKU fields. Specifically, every SKU value is one of potentially many SKU values associated with a particular PRODUCT_NAME value. Every PRODUCT_NAME value is one of potentially many PRODUCT_NAME values associated with a particular PRODUCT_TYPE value. Every PRODUCT_TYPE value is one of potentially many PRODUCT_TYPE values associated with a particular MANUFACTURER value.

The "product category hierarchy" implicit in the values contained in table 500 is illustrated in FIG. 6. Referring to FIG. 6, each of the four hierarchically related fields of table 500 correspond to a level of hierarchy, where the MANUFACTURER field is the highest level (just below the root node), and the SKU field is the lowest level (representing the "leaf" nodes).

The enumeration approach may be used to capture the parent-child relationships when pre-existing hierarchies are defined in this manner. Specifically, the one or more columns containing the enumerated values that are involved in the hierarchy are used to generate a "collection node". Each value of each of the columns is a new "collection" at a particular depth in the hierarchy, and each column is used for a new level of the hierarchy. The last column given is treated as a file, not a collection.

Specifically, the nodes in the first level of the hierarchy may be established by performing a DISTINCT operation on the values contained in the MANUFACTURER column. The nodes in the second level of the hierarchy may be established by concatenating the MANUFACTURE and PRODUCT_TYPE values for all rows, and performing a DISTINCT operation on the resulting concatenated values. Similarly, the nodes in the third level of the hierarchy may be established by concatenating the MANUFACTURE, PRODUCT_TYPE, and PRODUCT_NAME values for all rows, and performing a DISTINCT operation on the resulting concatenated values. Finally, the leaf level of the hierarchy may be established by concatenating the MANUFACTURE, PRODUCT_TYPE, PRODUCT_NAME and SKU values for all rows, and performing a DISTINCT operation on the resulting concatenated values.

The concatenation operations described above are performed under the assumption that the names of the nodes at a given level are not constrained to be unique. For example, it assumes that a product under the PRODUCT_TYPE value PD1 can have the same name as a product under the PRODUCT_TYPE value PD2. However, if the PRODUCT_NAME column is constrained to be unique, then the nodes at the PRODUCT_NAME level of the hierarchy may be obtained by simply performing a DISTINCT operation on the values in the PRODUCT_NAME column.

It should be noted that the relative pathname of each leaf node is indicated by the values stored in the relational table row that holds the resource for that node. For example, row 502 (FIG. 5) holds the resource for node 602 (FIG. 6) and the values in the columns of row 502 correspond to the relative pathname /MF1/PD2/PN3/SKU3. Consequently, when finding the children of a particular node (such as MF1/PD2) the

columns used for that node are implicitly used to select children (so the SQL used to find children of MF1/PD2 would be 'SELECT FROM PRODUCT_INFO WHERE manufacturer='MF1' and 'product_type' = 'PD2').

MAINTAINING CONSISTENCY

As mentioned above, after a pre-existing hierarchy has been captured, there exist two independent sets of data that reflect the hierarchy: the external hierarchy definition and the internal hierarchy definition. To maintain a consistent reflection of the hierarchy, the internal hierarchy definition must be modified in response to changes made to the external hierarchy definition, and the external hierarchy definition must be modified in response to changes made to the internal hierarchy definition.

For example, once captured, the management hierarchy associated with table 450 is reflected in (1) data stored in the hierarchy structures and (2) the relationship between the data in the EMP_ID column of table 450 and the data in the MANAGER column of table 450. Consequently, changes to the management hierarchy may be accomplished by either (1) changing a value in the EMP_ID or MANAGER columns of table 450 or (2) changing values in the hierarchy structures. To maintain consistency between the data in the hierarchy structures and the data in the EMP_ID and MANAGER columns, mechanisms are established to cause modifications made to one set of data to automatically cause corresponding modifications to the other set of data.

According to one embodiment, table 450 is altered such that certain checks (for example security checks) that are done as a part of the navigation of resource table 402 are done for traversing the table 450 too. Data Manipulation Language (DML) triggers are created on table 450 so that an insert, update or delete operation on the table 450 gets reflected in the data in the hierarchy structures. For example, if a row from table 450 is deleted, then a check is made to ensure that the row being deleted will not lead to any "dangling rows". For example, the node, within the management hierarchy, that is associated row 420 is the parent of the node associated with row 422. Consequently, the deletion of row 420 would sever row 422 from the management hierarchy, causing row 422 to "dangle". Because a dangling row is no longer part of the hierarchy, the deletion of row 422 would cause the deletion, from the hierarchy structures, of data relating to both row 420 and row 422.

Similarly, when a row is inserted into the table 450, the information in the hierarchy structures will be updated to add an entry for the new path that is created. Note that it is possible that when a new row is inserted, it adds a link that allows access to a

subtree that was not accessible earlier. For example, assume that after row 420 was deleted, row 420 is inserted once again. By inserting row 420, row 422 is implicitly grafted back into the management hierarchy. Consequently, the insertion of row 420 causes data in the hierarchy structures to be updated to include the nodes and paths for both row 420 and row 422.

The exact operations that are performed to maintain consistency between the internal hierarchy definition and the external hierarchy definition will vary based on a variety of factors. One factor that affects what must be done to maintain consistency is the nature of the external hierarchy definition. For example, the management hierarchy that applies to table 450 is based on an SQL relation, while the product category hierarchy that applies to table 500 is based on an enumeration. When a row associated with a leaf node in the management hierarchy is deleted, typically only the row in the resource table that corresponds to that node should also be deleted. However, if a row associated with a leaf node of the product category hierarchy is deleted, then (1) the row in the resource table that corresponds to that node is deleted, and (2) it is determined whether the parent of the deleted leaf node has any remaining children. If the parent of the deleted leaf node does not have any remaining children, then the row in the resource table that corresponds to the parent node is deleted. If the row that corresponds to the parent node is deleted, then it is determined whether the parent of the parent node has any remaining children. If the parent of the parent node does not have any remaining children, then the row of the resource table associated with the parent of the parent node is also deleted. This process is repeated to remove all childless non-leaf nodes.

Similarly, for drop and truncate operations, which involve deleting all rows in the table, operations for maintaining consistency, similar to those described above for the delete operation, are performed.

In the above examples, DML triggers perform modifications on data in the hierarchy structures based on insert, update and delete operations performed on the content structures. Such DML triggers may be implemented, for example, as "INSTEAD OF" triggers. However, to ensure consistency, triggers must also be created to ensure that insert, update and delete operations on the data in the hierarchy structures result in corresponding operations on the data in the content structures.

HARDWARE OVERVIEW

Figure 7 is a block diagram that illustrates a computer system 700 upon which an embodiment of the invention may be implemented. Computer system 700 includes a bus 702 or other communication mechanism for communicating information, and a processor 704 coupled with bus 702 for processing information. Computer system 700 also includes a main memory 706, such as a random access memory (RAM) or other dynamic storage device, coupled to bus 702 for storing information and instructions to be executed by processor 704. Main memory 706 also may be used for storing temporary variables or other intermediate information during execution of instructions to be executed by processor 704. Computer system 700 further includes a read only memory (ROM) 708 or other static storage device coupled to bus 702 for storing static information and instructions for processor 704. A storage device 710, such as a magnetic disk or optical disk, is provided and coupled to bus 702 for storing information and instructions.

Computer system 700 may be coupled via bus 702 to a display 712, such as a cathode ray tube (CRT), for displaying information to a computer user. An input device 714, including alphanumeric and other keys, is coupled to bus 702 for communicating information and command selections to processor 704. Another type of user input device is cursor control 716, such as a mouse, a trackball, or cursor direction keys for communicating direction information and command selections to processor 704 and for controlling cursor movement on display 712. This input device typically has two degrees of freedom in two axes, a first axis (e.g., x) and a second axis (e.g., y), that allows the device to specify positions in a plane.

The invention is related to the use of computer system 700 for implementing the techniques described herein. According to one embodiment of the invention, those techniques are performed by computer system 700 in response to processor 704 executing one or more sequences of one or more instructions contained in main memory 706. Such instructions may be read into main memory 706 from another computer-readable medium, such as storage device 710. Execution of the sequences of instructions contained in main memory 706 causes processor 704 to perform the process steps described herein. In alternative embodiments, hard-wired circuitry may be used in place of or in combination with software instructions to implement the invention. Thus, embodiments of the invention are not limited to any specific combination of hardware circuitry and software.

The term "computer-readable medium" as used herein refers to any medium that participates in providing instructions to processor 704 for execution. Such a medium may take many forms, including but not limited to, non-volatile media, volatile media, and transmission media. Non-volatile media includes, for example, optical or magnetic disks, such as storage device 710. Volatile media includes dynamic memory, such as main memory 706. Transmission media includes coaxial cables, copper wire and fiber optics, including the wires that comprise bus 702. Transmission media can also take the form of acoustic or light waves, such as those generated during radio-wave and infra-red data communications.

Common forms of computer-readable media include, for example, a floppy disk, a flexible disk, hard disk, magnetic tape, or any other magnetic medium, a CD-ROM, any other optical medium, punchcards, papertape, any other physical medium with patterns of holes, a RAM, a PROM, and EPROM, a FLASH-EPROM, any other memory chip or cartridge, a carrier wave as described hereinafter, or any other medium from which a computer can read.

Various forms of computer readable media may be involved in carrying one or more sequences of one or more instructions to processor 704 for execution. For example, the instructions may initially be carried on a magnetic disk of a remote computer. The remote computer can load the instructions into its dynamic memory and send the instructions over a telephone line using a modem. A modem local to computer system 700 can receive the data on the telephone line and use an infra-red transmitter to convert the data to an infra-red signal. An infra-red detector can receive the data carried in the infra-red signal and appropriate circuitry can place the data on bus 702. Bus 702 carries the data to main memory 706, from which processor 704 retrieves and executes the instructions. The instructions received by main memory 706 may optionally be stored on storage device 710 either before or after execution by processor 704.

Computer system 700 also includes a communication interface 718 coupled to bus 702. Communication interface 718 provides a two-way data communication coupling to a network link 720 that is connected to a local network 722. For example, communication interface 718 may be an integrated services digital network (ISDN) card or a modem to provide a data communication connection to a corresponding type of telephone line. As another example, communication interface 718 may be a local area network (LAN) card to provide a data communication connection to a compatible LAN. Wireless links may also be implemented. In any such implementation, communication

interface 718 sends and receives electrical, electromagnetic or optical signals that carry digital data streams representing various types of information.

Network link 720 typically provides data communication through one or more networks to other data devices. For example, network link 720 may provide a connection through local network 722 to a host computer 724 or to data equipment operated by an Internet Service Provider (ISP) 726. ISP 726 in turn provides data communication services through the world wide packet data communication network now commonly referred to as the "Internet" 728. Local network 722 and Internet 728 both use electrical, electromagnetic or optical signals that carry digital data streams. The signals through the various networks and the signals on network link 720 and through communication interface 718, which carry the digital data to and from computer system 700, are exemplary forms of carrier waves transporting the information.

Computer system 700 can send messages and receive data, including program code, through the network(s), network link 720 and communication interface 718. In the Internet example, a server 730 might transmit a requested code for an application program through Internet 728, ISP 726, local network 722 and communication interface 718.

The received code may be executed by processor 704 as it is received, and/or stored in storage device 710, or other non-volatile storage for later execution. In this manner, computer system 700 may obtain application code in the form of a carrier wave.

In the foregoing specification, the invention has been described with reference to specific embodiments thereof. It will, however, be evident that various modifications and changes may be made thereto without departing from the broader spirit and scope of the invention. The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense.

CLAIMS

What is claimed is:

1. A method for managing data in a database system, comprising the steps of:
 - receiving data that identifies one or more relational tables that store resources that belong to an information hierarchy, wherein certain data within said one or more relational tables serves as an external hierarchy definition that establishes hierarchical relationships between said resources;
 - capturing the hierarchical relationships within said hierarchy by performing the steps of
 - executing one or more queries against the one or more relational tables to retrieve said certain data; and
 - based on said certain data, storing information that captures said information hierarchy, as an internal hierarchy definition, in one or more hierarchy structures separate from said one or more relational tables.
2. The method of Claim 1 wherein the step of storing information that captures said information hierarchy includes storing, within a resource table, a row for each node in said information hierarchy.
3. The method of Claim 1 wherein the step of storing information that captures said information hierarchy includes storing, in one or more hierarchy structures, information that indicates parent-child relationships between said resources.
4. The method of Claim 1 wherein:
 - the resources in the information hierarchy correspond to rows of a relational table;
 - and
 - parent-child relationships between the resources are indicated by values stored in two or more columns of the relational table.
5. The method of Claim 4 wherein the step of executing one or more queries includes executing a query that specifies a CONNECT BY relationship between said two or more columns.

6. The method of Claim 5 wherein:
the two or more columns include a first column that contains data indicating a child in a relationship and a second column that contains data indicating a parent in the relationship; and
the step of executing a query includes executing a query that contains an expression in the form: START WITH x IS NULL CONNECT BY PRIOR x = y;
wherein x represents the second column and y represents the first column.
7. The method of Claim 1 wherein:
the resources associated with leaf nodes in the information hierarchy correspond to rows of a relational table; and
values associated with non-leaf nodes that, within the information hierarchy, are above a given leaf node are stored in the row that corresponds to the given leaf node.
8. The method of Claim 7 wherein:
the values associated with non-leaf nodes are stored in one or more columns of said relational table; and
the step of storing information that captures said information hierarchy includes storing, within a resource table, a row for each distinct value, within said one or more columns, from rows that correspond to leaf nodes in the information hierarchy.
9. The method of Claim 1 further comprising the step of establishing triggers to update the data that constitutes said internal hierarchy definition in response to changes made to data that constitutes the external hierarchy definition.
10. The method of Claim 1 further comprising the step of establishing triggers to update the data that constitutes said external hierarchy definition in response to changes made to data that constitutes the internal hierarchy definition.

11. The method of Claim 1 further comprising the step of automatically maintaining consistency between said external hierarchy definition and said internal hierarchy definition.
12. The method of Claim 11 wherein the step of automatically maintaining consistency includes the steps of:
 - detecting when an operation on said one or more tables deletes a resource associated with a node in the information hierarchy;
 - in response to detecting said operation, automatically performing the steps of deleting, in a resource table, a row that corresponds to said node;
 - determining whether removal of said node from said information hierarchy causes one or more other nodes to dangle; and
 - if removal of said node from said information hierarchy causes one or more other nodes to dangle, then deleting, in the resource table, rows that correspond to said one or more other nodes.
13. The method of Claim 11 wherein the step of automatically maintaining consistency includes the steps of:
 - detecting when an operation on said one or more tables deletes a resource associated with a leaf node in the information hierarchy;
 - in response to detecting said operation, automatically performing the steps of deleting, in a resource table, a row that corresponds to said node;
 - determining whether removal of said leaf node from said information hierarchy leaves a parent node of said leaf node with no remaining child nodes; and
 - if removal of said leaf node from said information hierarchy leaves a parent node of said leaf node with no remaining child nodes, then deleting, in the resource table, a row that corresponds to said parent node.
14. The method of Claim 11 wherein the step of automatically maintaining consistency includes the steps of:
 - detecting when an operation on said one or more tables inserts a resource associated with a node in the information hierarchy;

in response to detecting said operation, automatically performing the steps of inserting, in a resource table, a row that corresponds to said node; determining whether addition of said node from said information hierarchy causes the addition of one or more child nodes of said node to said information hierarchy; and if addition of said node from said information hierarchy causes the addition of one or more child nodes of said node to said information hierarchy, then inserting, in the resource table, rows that correspond to said one or more child nodes.

15. A computer-readable medium carrying instructions for managing data in a database system, the instructions comprising instructions for performing the steps of:
 - receiving data that identifies one or more relational tables that store resources that belong to an information hierarchy, wherein certain data within said one or more relational tables serves as an external hierarchy definition that establishes hierarchical relationships between said resources;
 - capturing the hierarchical relationships within said hierarchy by performing the steps of
 - executing one or more queries against the one or more relational tables to retrieve said certain data; and
 - based on said certain data, storing information that captures said information hierarchy, as an internal hierarchy definition, in one or more hierarchy structures separate from said one or more relational tables.

16. The computer-readable medium of Claim 15 wherein the step of storing information that captures said information hierarchy includes storing, within a resource table, a row for each node in said information hierarchy.

17. The computer-readable medium of Claim 15 wherein the step of storing information that captures said information hierarchy includes storing, in one or more hierarchy structures, information that indicates parent-child relationships between said resources.

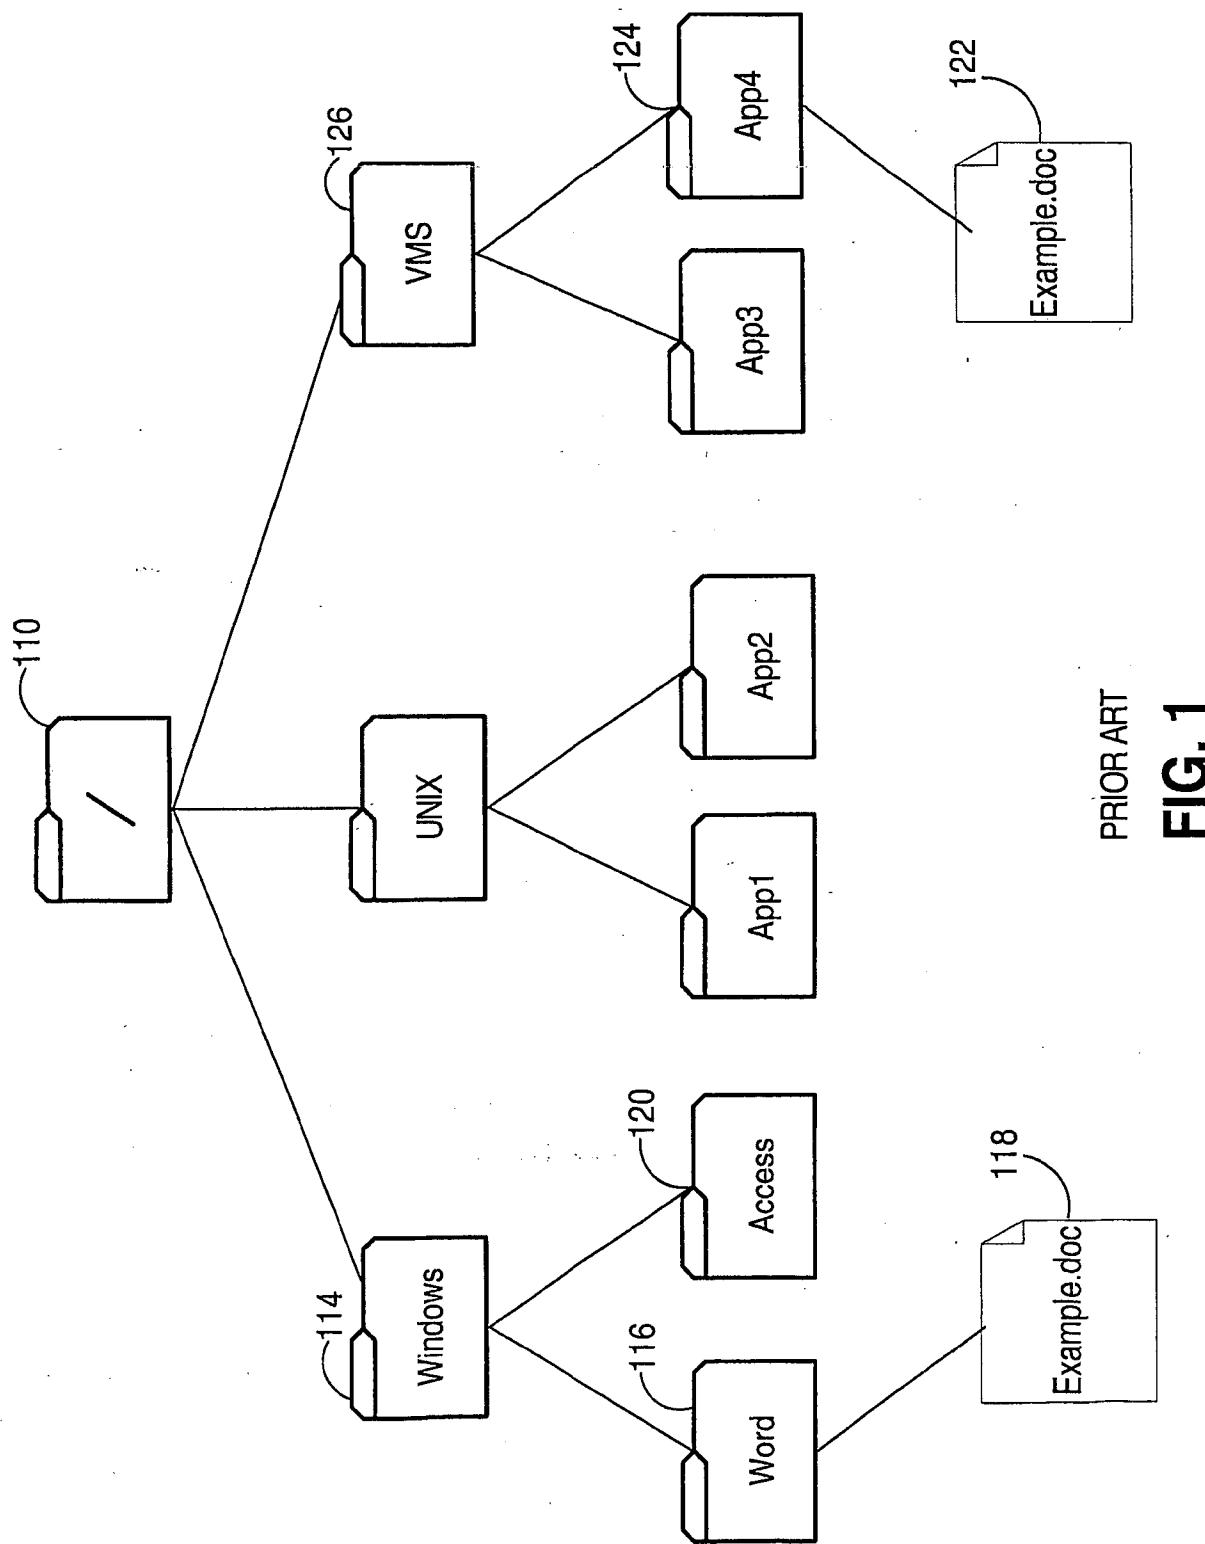
18. The computer-readable medium of Claim 15 wherein:
the resources in the information hierarchy correspond to rows of a relational table;
and
parent-child relationships between the resources are indicated by values stored in
two or more columns of the relational table.
19. The computer-readable medium of Claim 18 wherein the step of executing one or
more queries includes executing a query that specifies a CONNECT BY relationship
between said two or more columns.
20. The computer-readable medium of Claim 19 wherein:
the two or more columns include a first column that contains data indicating a
child in a relationship and a second column that contains data indicating a
parent in the relationship; and
the step of executing a query includes executing a query that contains an
expression in the form: START WITH x IS NULL CONNECT BY
PRIOR x = y;
wherein x represents the second column and y represents the first column.
21. The computer-readable medium of Claim 15 wherein:
the resources associated with leaf nodes in the information hierarchy correspond
to rows of a relational table; and
values associated with non-leaf nodes that, within the information hierarchy, are
above a given leaf node are stored in the row that corresponds to the given
leaf node.
22. The computer-readable medium of Claim 21 wherein:
the values associated with non-leaf nodes are stored in one or more columns of
said relational table; and
the step of storing information that captures said information hierarchy includes
storing, within a resource table, a row for each distinct value, within said
one or more columns, from rows that correspond to leaf nodes in the
information hierarchy.

23. The computer-readable medium of Claim 15 further comprising instructions for performing the step of establishing triggers to update the data that constitutes said internal hierarchy definition in response to changes made to data that constitutes the external hierarchy definition.
24. The computer-readable medium of Claim 15 further comprising instructions for performing the step of establishing triggers to update the data that constitutes said external hierarchy definition in response to changes made to data that constitutes the internal hierarchy definition.
25. The computer-readable medium of Claim 15 further comprising instructions for performing the step of automatically maintaining consistency between said external hierarchy definition and said internal hierarchy definition.
26. The computer-readable medium of Claim 25 wherein the step of automatically maintaining consistency includes the steps of:
- detecting when an operation on said one or more tables deletes a resource associated with a node in the information hierarchy;
- in response to detecting said operation, automatically performing the steps of deleting, in a resource table, a row that corresponds to said node;
- determining whether removal of said node from said information hierarchy causes one or more other nodes to dangle; and
- if removal of said node from said information hierarchy causes one or more other nodes to dangle, then deleting, in the resource table, rows that correspond to said one or more other nodes.
27. The computer-readable medium of Claim 25 wherein the step of automatically maintaining consistency includes the steps of:
- detecting when an operation on said one or more tables deletes a resource associated with a leaf node in the information hierarchy;
- in response to detecting said operation, automatically performing the steps of deleting, in a resource table, a row that corresponds to said node;

determining whether removal of said leaf node from said information hierarchy leaves a parent node of said leaf node with no remaining child nodes; and
if removal of said leaf node from said information hierarchy leaves a parent node of said leaf node with no remaining child nodes, then deleting, in the resource table, a row that corresponds to said parent node.

28. The computer-readable medium of Claim 25 wherein the step of automatically maintaining consistency includes the steps of:
 - detecting when an operation on said one or more tables inserts a resource associated with a node in the information hierarchy;
 - in response to detecting said operation, automatically performing the steps of inserting, in a resource table, a row that corresponds to said node;
 - determining whether addition of said node from said information hierarchy causes the addition of one or more child nodes of said node to said information hierarchy; and
 - if addition of said node from said information hierarchy causes the addition of one or more child nodes of said node to said information hierarchy, then inserting, in the resource table, rows that correspond to said one or more child nodes.

1/7



2/7

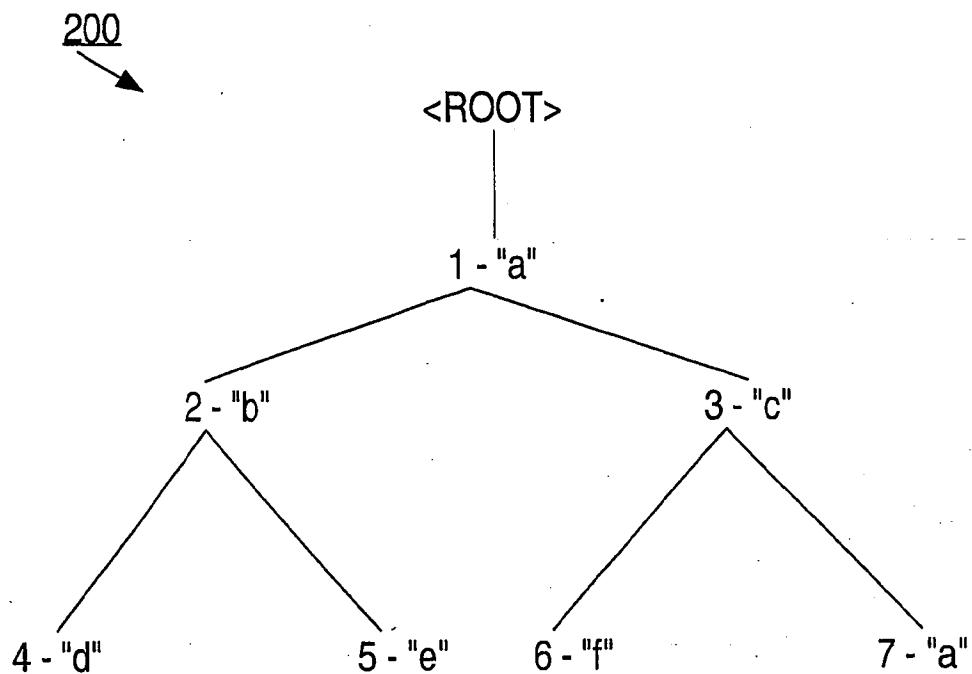


FIG. 2

3/7

FIG. 3

The top table is a node database:

NODE	NAME	DATA
1	a	D1
2	b	D2
3	c	D3
4	d	D4
5	e	D5
6	f	D6
7	a	D7

Arrows labeled 302, 304, 308, and 306 point to the first row, second column, third column, and the last row of the top table respectively.

The bottom table is a parent-child relationship table:

PARENT	CHILD
<ROOT>	1
1	2
1	3
2	4
2	5
3	6
3	7

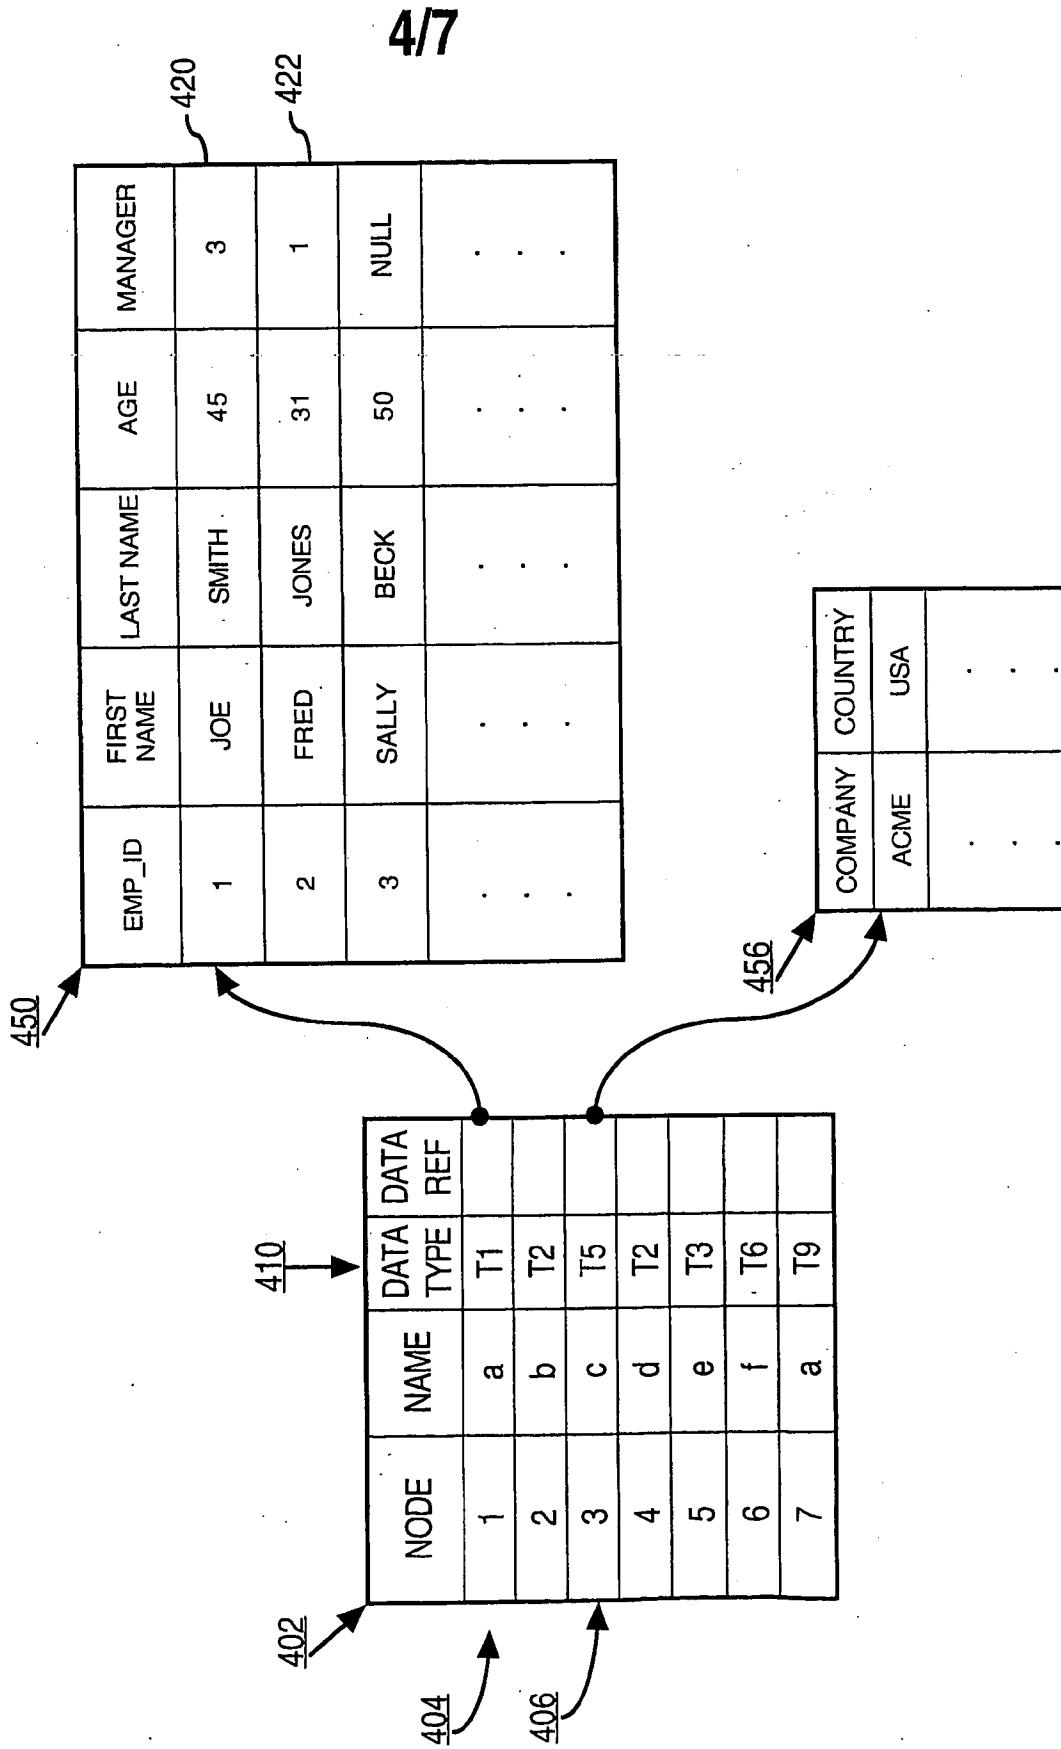
Arrows labeled 350, 378, 380, 352, 354, 356, 358, 360, 362, and 364 point to the first column, second column, and specific rows of the bottom table.

The bottom table is a parent-child relationship table:

PARENT	CHILD
<ROOT>	1
1	2
1	3
2	4
2	5
3	6
3	7

Arrows labeled 350, 378, 380, 352, 354, 356, 358, 360, 362, and 364 point to the first column, second column, and specific rows of the bottom table.

FIG. 4



5/7

MANUFACTURER	PRODUCT_TYPE	PRODUCT_NAME	SKU
MF1	PD1	PN1	SKU1
MF1	PD1	PN2	SKU2
MF1	PD2	PN3	SKU3
MF2	PD3	PN4	SKU4
MF2	PD3	PN4	SKU5
.	.	.	.
.	.	.	.

FIG. 5

6/7

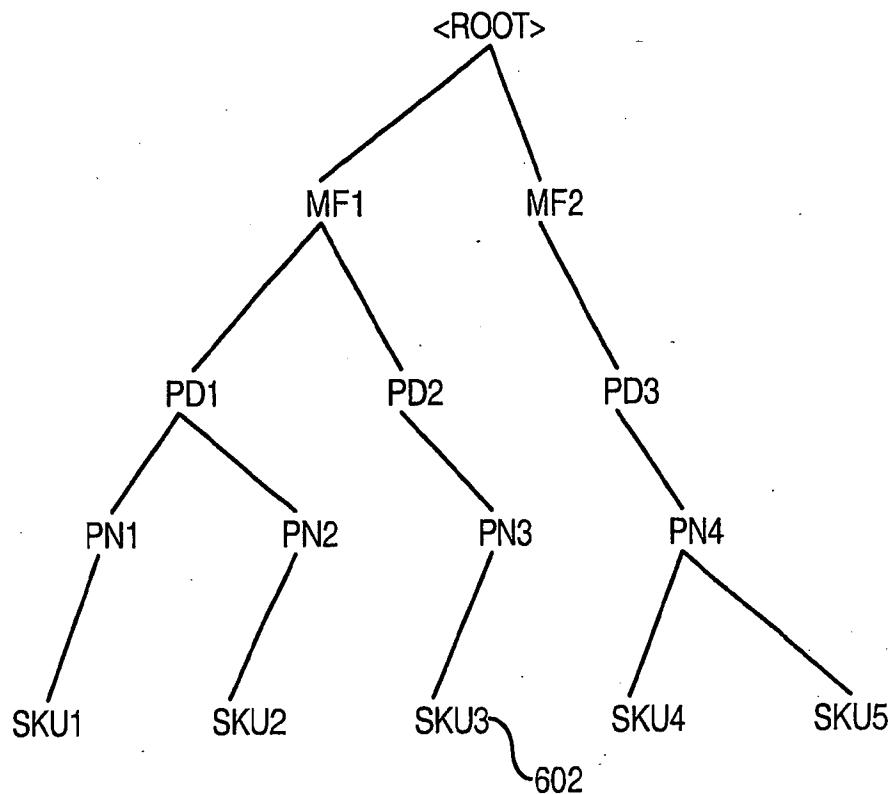


FIG. 6

7/7

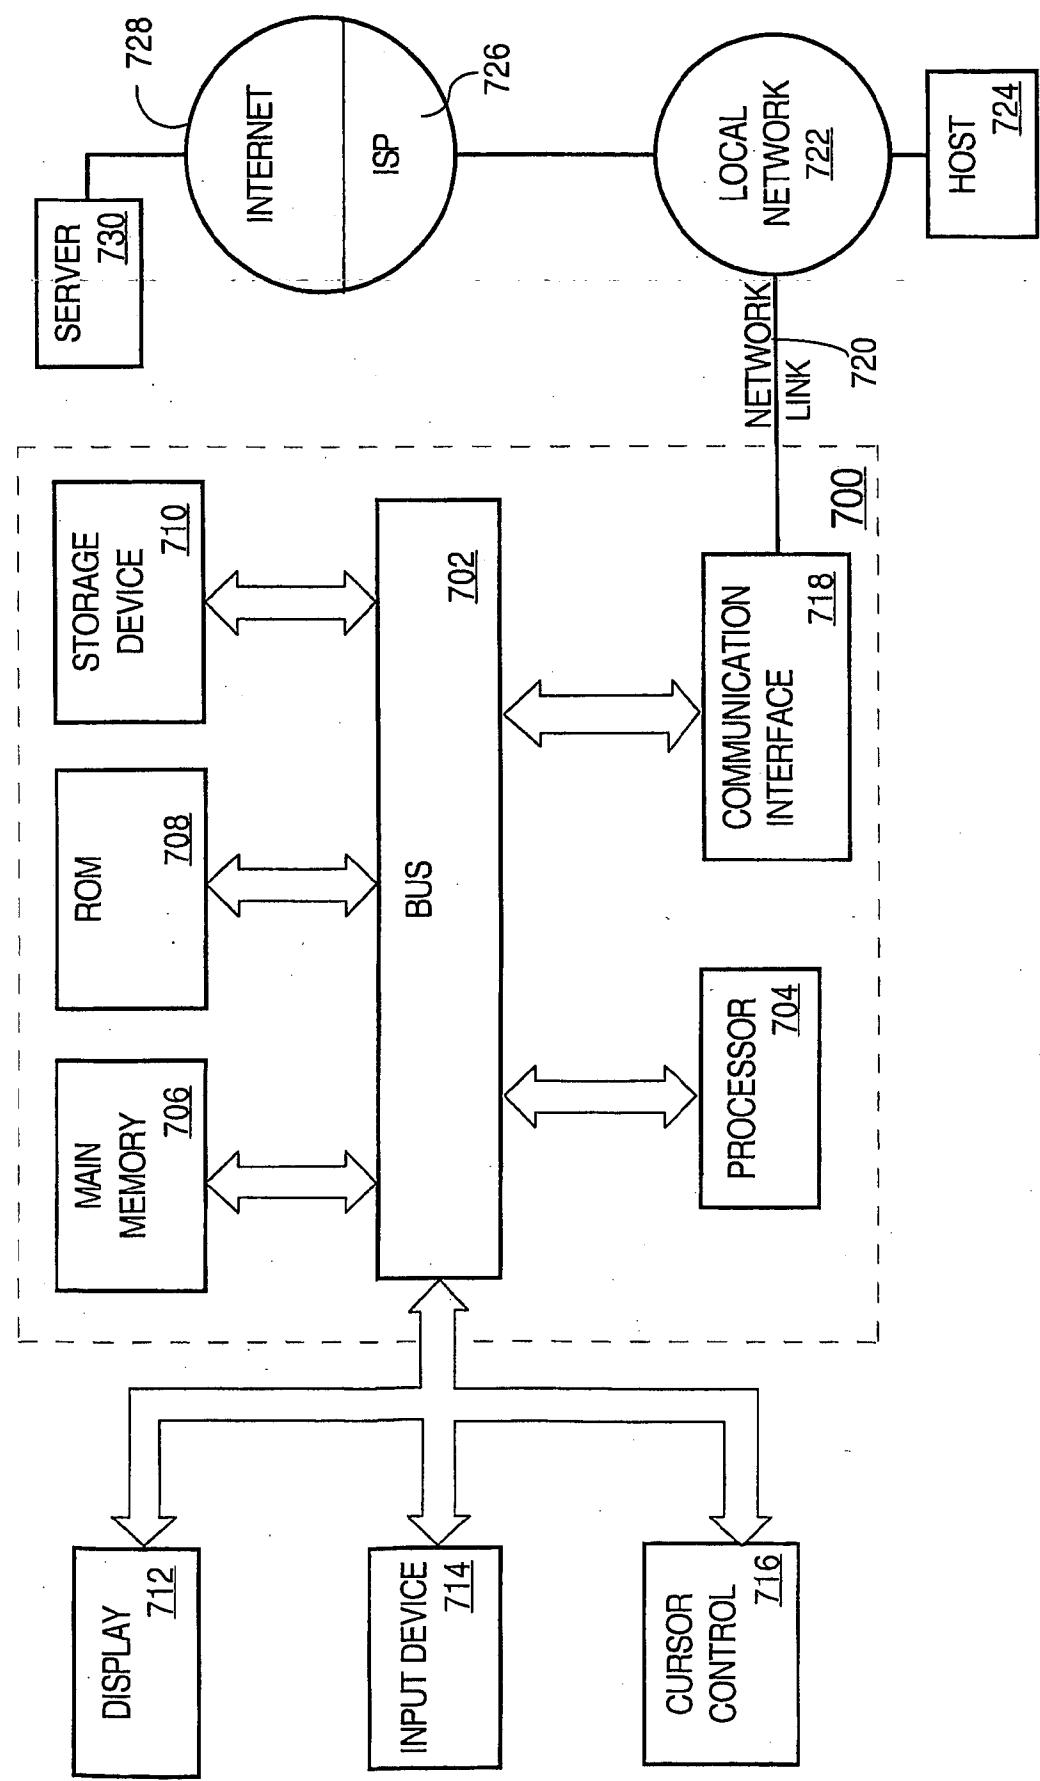


FIG. 7